Journal of Astronomical Instrumentation
© World Scientific Publishing Company

arXiv:1408.6529v1 [astro-ph.IM] 27 Aug 2014

Characterization of a commercial, front-illuminated interline transfer CCD camera for use as a guide camera on a balloon-borne telescope

Paul Clark[1], Richard Massey[1,2], Herrick L. Chang[3], Mathew Galloway[4], Holger Israel[2], Laura L. Jones[3], Lun Li[4], Milan Mandic[3], Tim Morris[1], Barth Netterfield[4], John Peacock[5], Ray Sharples[1] and Sara Susca[3]
[1]*Centre for Advanced Instrumentation, University of Durham, Joseph Swan Road, Netpark, Sedgefield TS21 3FB, United Kingdom*
[2]*Institute for Computational Cosmology, University of Durham, South Road, Durham DH1 3LE, United Kingdom*
[3]*Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, United States*
[4]*University of Toronto, Department of Physics, 60 St. George St., Toronto, Ontario M5S 1A7, Canada*
[5]*Institute for Astronomy, University of Edinburgh, Royal Observatory Blackford Hill, Edinburgh, EH9 3HJ, United Kingdom*

Received June 27, 2014; Revised August 26, 2014; Accepted August 27, 2014.

We report results obtained during the characterisation of a commercial front-illuminated progressive scan interline transfer CCD camera. We demonstrate that the unmodified camera operates successfully in temperature and pressure conditions (−40°C, 4 mBar) representative of a high altitude balloon mission. We further demonstrate that the centroid of a well-sampled star can be determined to better than 2% of a pixel, even though the CCD is equipped with a microlens array. This device has been selected for use in a closed-loop star-guiding and tip-tilt correction system in the BIT-STABLE balloon mission.

Keywords: Instrumentation, CCD characterization, sub-orbital balloon flight, pointing stability

1. Introduction

High altitude balloons offer a revolutionary environment to conduct optical imaging and spectroscopy for astronomy and Earth observation. From an altitude of 35 km (above 99% of the Earth's atmosphere), image quality potentially equals that of satellites in near-Earth orbit, but at a fraction (∼1%) of the cost and with the opportunity to repair or upgrade hardware between flights (Rhodes et al. 2012). Furthermore, the recent development of Ultra Long Duration, super pressure balloons (ULDBs) allows large surveys during flights that last several months.

Balloon missions face limited budgets on cost and complexity (and size and mass). Commercial, off-the-shelf hardware can be adopted to stay within these budgets, but only if its performance meets requirements. In this paper, we describe tests of an Imperx Bobcat IGV-B2320-M to (i) operate in a flight-like temperature and pressure environment (-40°C, 4 mBar) and (ii) perform science-quality imaging. This camera was one of several candidates considered for the star-tracking guide camera of the BIT-STABLE project, a collaboration between NASA/JPL, the University of Toronto, and a UK consortium (Durham University and the University of Edinburgh) to fly an optical telescope on a high altitude balloon.

The Imperx camera is available with an ethernet interface that simplifies interfacing to the flight computer and makes a flight-compatible framegrabber unnecessary. Its internal electronics have already been ruggedized for use in industrial settings, for example with a thermally-optimised (fanless) design and internal conformal coating. We use an environmental chamber to recreate the temperature and pressure cycle of a typical balloon mission, and explicitly check that the camera will both survive and continue to operate.

The imaging detector is a Truesense KAI-04050 front-illuminated interline transfer CCD, with 2352×1768 $5.5\,\mu$m (square) pixels and 12-bit readout. Interline transfer CCDs have a non-responsive area in each pixel, which is used for fast readout. This can usefully mimic an electronic shutter, reducing complexity. However, a lenslet array is then needed to focus incident light onto the responsive area in each pixel to restore QE, and distortion through these optical elements may spuriously move the apparent position of a guide star. Using a precision translation stage, we move a spot, representative of the BIT-STABLE PSF, within several

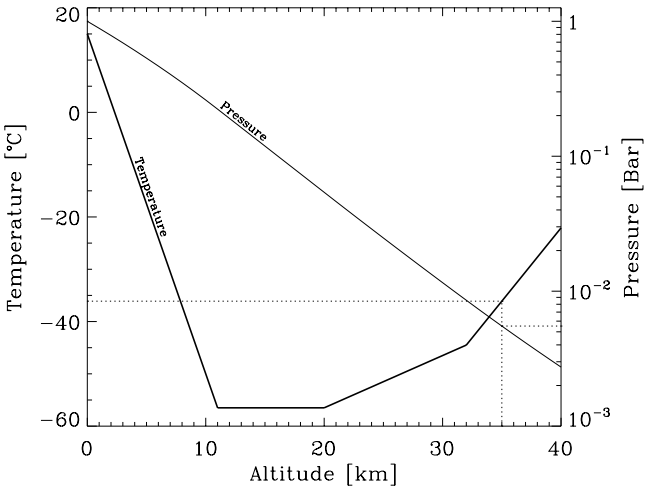

Fig. 1. Environmental conditions as a function of altitude, in the 1976 U.S. Standard Atmosphere model.

of the device's pixels. We then explicitly compare the known position of the spot to the position of the pixel in which light is actually registered. During testing, we also noticed anomalies in images obtained in Windowed Readout Mode that could also affect position measurements.

This paper is organized as follows. In section 2, we describe camera operational tests in an environmental chamber that replicates flight conditions. In section 3, we describe measurements of sub-pixel variations in detector response. In section 4, we describe measurements of anomalies in windowed captures, investigate their possible origin, and develop a mitigation strategy. We conclude in section 5.

2. Operation under flight conditions

2.1. *The problem*

The BIT-STABLE balloon mission will operate at an altitude of 35 km. As seen from the U.S. Standard Atmosphere [see figure 1; U.S. Government Printing Office, 1976], the ambient pressure and temperature at that altitude are approximately −40°C and 4 mBar (0.004 atm). The pressure falls exponentially during ascent, while the temperature profile passes through a minimum of −56°C at an altitude of 11–20 km before rising again towards float altitude. The ascent to 35 km is expected to take approximately 100 minutes (based on balloon flight data from BLAST; Truch *et al.* 2009), so the balloon hardware will only encounter ambient temperatures as low as −56°C for a relatively short time. Since the telescope structure and gondola provide significant thermal mass and some heating, the BIT-STABLE mission has defined requirements of −40°C survivability temperature and −30°C minimum operating temperature for all flight hardware.

As the air pressure decreases with altitude, convective cooling also decreases, perversely resulting in the flight hardware being able to overheat rather than becoming too cold. BLAST used air-filled pressure vessels to maintain convective cooling of flight electronics [Pascale *et al.*, 2008]; BIT-STABLE's financial and weight budgets (the camera must be moved to track the sky) make pressure vessels undesirable. It is therefore essential to test candidate electronics, including the guide camera, under representative flight conditions to ensure they can operate reliably in the low temperature, low pressure, low convective cooling environment.

The Imperx Bobcat range of cameras are available with an internal conformal coating and have a thermally optimised, fan-less design allowing them to operate over an extended temperature range of −40°C to +85°C (in air). Operation at high altitude should be possible so long as the ∼5 W internal power dissipation doesn't raise the internal temperature to unsafe levels, or the low pressure cause failure of (e.g.) electrolytic can capacitors. The project team at Durham University therefore subjected an IGV-B2320 to a series of environmental tests prior to shipping to JPL. Initial tests concentrated on establishing the level of self-heating in the camera at reduced pressure. Later tests established the camera's ability to survive and operate in the low pressure and temperature environment.

2.2. *Experimental setup*

2.2.1. *Experimental configuration*

To replicate flight conditions in the laboratory, a bespoke cryogenic test chamber was used, as illustrated in figures 2–4. The test chamber contains a 50 cm diameter cold bench cooled by two CTI Cryo-Torr cold heads. A cylindrical radiation shield and multi-layer insulation (MLI) prevent the cold mass from being warmed by radiation from the room temperature vessel. During operation the cold bench stabilises at a temperature of approximately −240°C, much too cold for the project's needs. The Imperx camera was therefore mounted inside a secondary radiation shield, an aluminium alloy box, connected to the cold bench by thermal braids. The cross-sectional area of the cold braids between the box and the cold bench were tuned to ensure the box would be cooled to approximately the desired tem-

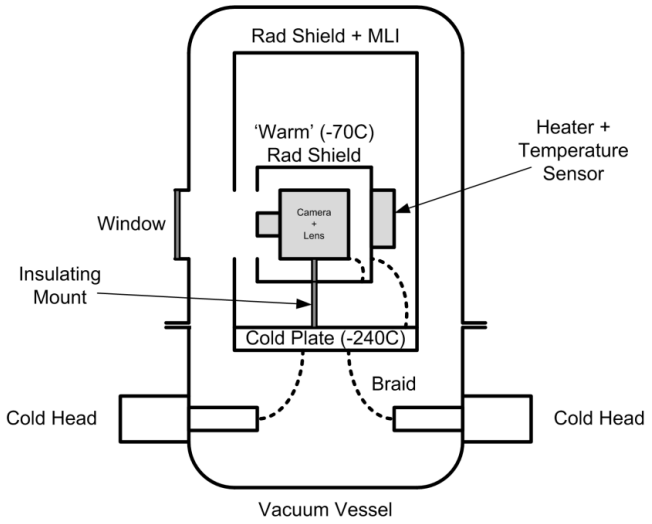

Fig. 2. Diagram of components inside the cryogenic chamber used to test camera survivability and operability at flight temperature and pressure.

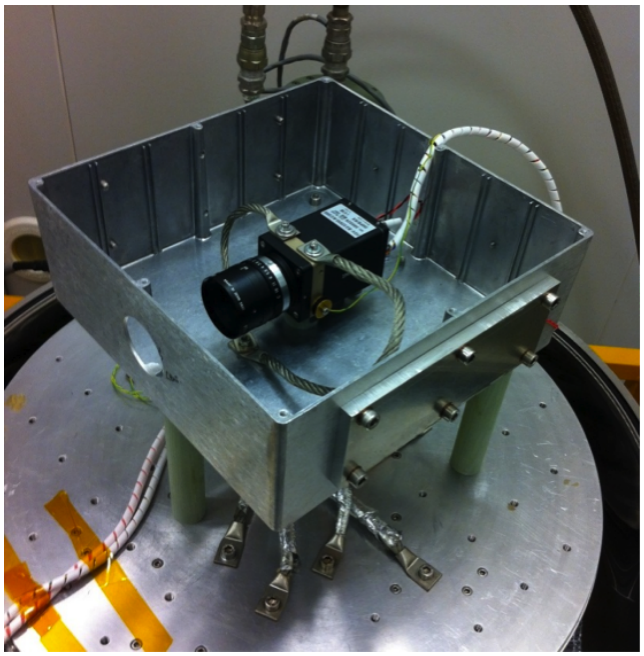

Fig. 3. Photograph of the camera inside the secondary radiation shield (see figure 2), with the lid removed. A pair of internal thermal braids connects the camera to the radiation shield; these were absent during initial runs to measure self-heating.

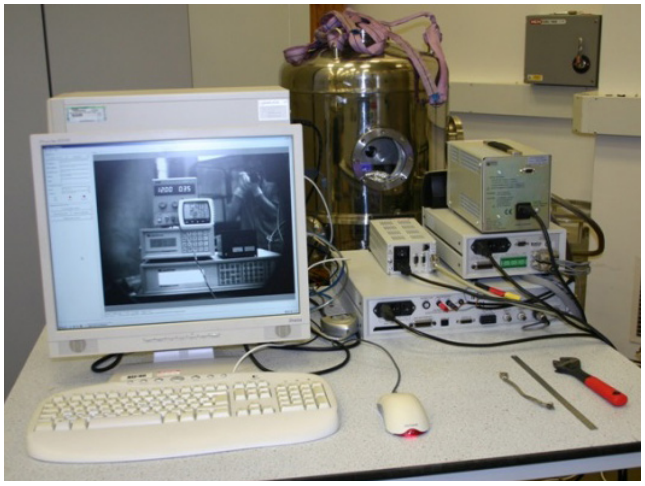

Fig. 4. Photograph of the closed cryogenic chamber (see figure 2) and support equipment. The monitor shows an image taken by the camera looking out from inside the chamber, showing the power supply and Lakeshore controller etc.

perature, then more finely regulated by two polyimide film heaters mounted on the box and controlled by an external Lakeshore temperature controller.

Initial tests were conducted with the camera isolated from the box, to determine the level of self-heating when operating in near-vacuum. For subsequent tests, small diameter cold braids (visible in figure 3) were installed between the camera and the box. By regulating the temperature of the box to approximately $-70°$C, the camera temperature could be lowered to the required $-30°$C (operating) or $-40°$C (non-operating, survival). Vacuum feed-through connectors allowed the camera power and ethernet connections to be brought to the outside. A hole was cut in the radiation shield to allow the camera to take images through the chamber window, of the Lakeshore controller, Lakeshore temperature monitor and Leybold vacuum gauge located adjacent to the chamber (figure 4). Temperatures were recorded via Lakeshore diode sensors. An analogue wallclock provided a useful time reference when the images taken periodically by the camera were turned into a time lapse movie sequence.

2.2.2. *Camera GigE interface*

On three separate computers (2 Windows, 1 Linux), the camera's Gigabit ethernet connection would occasionally produce dropped frames or partial images. We postulate that there is a certain dependence either on the ethernet drivers or ethernet hardware that is not fully compatible with the GigE Vision standard. By forcing the computers' connection back to 100 Mbps Full Duplex mode, this anomaly disappeared. This workaround will bottleneck the maximum frame rate for large windows or full frame captures. However, occasional full frame readout, or 50 Hz readout of windows smaller

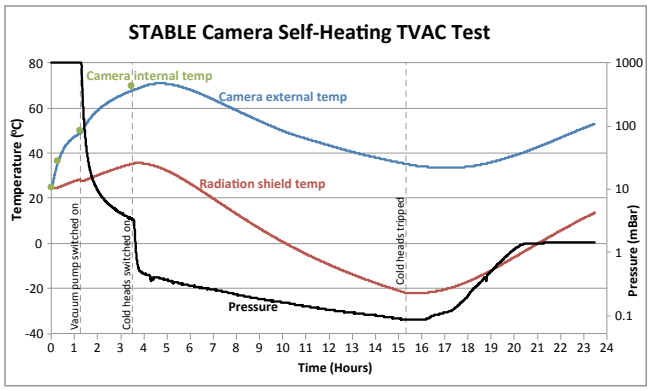

Fig. 5. Temperature and pressure profile during initial thermal cycle to measure self-heating (without cold braids connected to the camera).

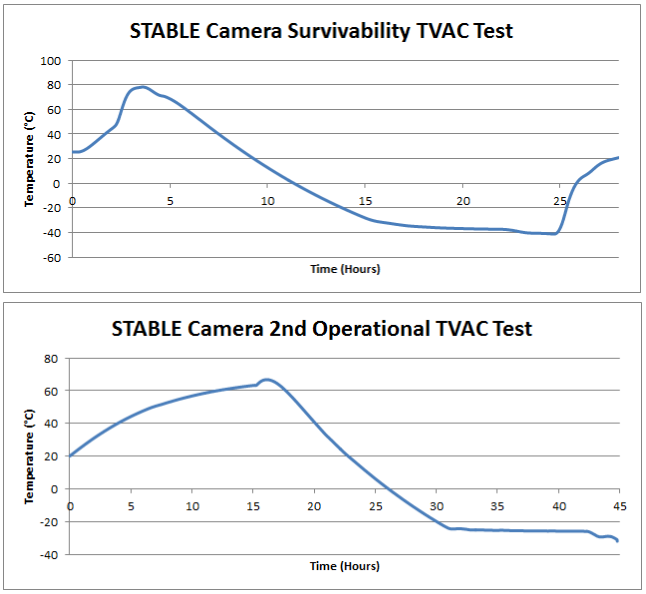

Fig. 6. Temperature profile during survivability and operational tests of the camera (with cold braids connected to the camera, enabling it to be sufficiently cooled).

than $\sim 256 \times 256$ pixels, remains possible within the 100 Mbps ethernet limit.

2.3. *Results*

2.3.1. *Self-heating*

To measure the level of self-heating in the camera, cold braids between the camera and the secondary radiation shield were omitted for an initial thermal cycle. The camera was operated continuously throughout the test, with full frame images recorded every minute (later to be turned into a time-lapse sequence), although power consumption of the camera did not vary measurably if images were taken at a higher frame rate. The recorded temperature and pressure profiles are shown in figure 5.

- T=0: the camera was powered up, in atmosphere. The temperature of the camera body, recorded via a Lakeshore temperature monitor, rose above the 22°C ambient to 50°C. The camera internal temperature, reported by the GEV software, was logged occasionally and observed to stabilise at approximately 2°C above the temperature of the camera body — indicating good thermal contact between the two.
- T=1.25 hours: the vacuum pump was switched on and the pressure inside the vacuum chamber was seen to fall. The vacuum pump, a large dry scroll pump, was connected to the chamber via 6 mm plastic hose, rather than conventional large bore vacuum hose, to deliberately reduce the decrease in pressure to a rate similar to what the camera will encounter during flight. The camera exhibits a corresponding increase in temperature as the convective cooling is reduced.
- T=3.4 hours: the cold heads were switched on when the camera reached a maximum temperature of 70°C. With the cold heads running, the pressure is seen to drop more rapidly as cryo pumping takes place. This is unavoidable and unfortunately takes the camera below the operational pressure of 4 mBar. The temperature of the camera was observed to fall, matching the fall in the temperature of the secondary radiation shield.
- T=15.3 hours: the cold heads were switched off causing the temperatures and pressure to rise accordingly.

The camera stabilises at ~ 45°C above ambient in vacuum without cold braids, due to ~ 5 W internal power dissipation. If the camera were mounted on an insulating mount on the BIT-STABLE telescope, we could therefore expect the camera temperature to remain above 0°C throughout the flight: reducing the risk of frosting during ascent or decent. However, a desire to minimise dark current in the CCD may require the camera to be thermally strapped to the cold mass of the telescope.

2.3.2. *Survivability and operational tests*

For the subsequent survivability and operational tests, cold braids were installed between the camera body and the inner radiation shield, as shown in figure 3. These allowed the camera to be cooled to the required −40°C survivability and −30°C operating temperatures. The survivability test required

the camera to be warmed to +70°C, dwell for one hour, then cooled to −40°C, dwell for one hour, and then return to ambient. The camera was not powered during this test. The operational test required the camera to be warmed to +60°C, dwell for one hour, then cooled to −30°C, dwell for one hour, and then returned to ambient. The camera was continuously operational throughout this test. The operational test was repeated a second time. Figure 6 shows the camera body temperature profiles during the survivability and second operational tests.

As the camera under test would be used for flight, it was necessary to avoid subjecting it to any unnecessary temperature cycles. This meant that setting the temperature of the secondary radiation shield to achieve the desired camera temperature involved a little educated guesswork. With repeated testing the target temperatures could be determined more accurately, but for the first time tests shown in figure 6, the temperature profiles are certainly non-optimal. During the survivability test for example, the secondary radiation shield temperature was initially set a little above +70°C. After 2.5 hours it became clear that the (non-operating) camera was going to take a long time to reach the required +70°C, so the radiation shield temperature was raised. This then resulted in the camera temperature overshooting and reaching close to +80°C. Having been above +70°C for the required one hour, cooling was started and the secondary radiation shield temperature was set to a little below −40°C. At T=22 hours, the following morning, the camera had not quite reached −40°C, so the radiation shield temperature was lowered. The camera then spent the required one hour at −40°C before the cooling was switched off and the camera warmed back to ambient. The profile for the operational test is similar with the radiation shield temperatures having to be nudged to achieve the target +60°C / −30°C (operating) camera temperature.

The camera passed the survivability and both operational tests, and continues to function well during integration into the telescope system at JPL.

3. Sub-pixel Photo Response Non-Uniformity

3.1. *The problem*

The architecture of an interline CCD is ideal for a balloon telescope's guide camera in several ways. At the end of an exposure, charge is stepped sideways by one pixel, into a column that is non-responsive to light but dedicated to charge transfer and readout. This architecture avoids the need for a mechanical shutter, which would otherwise be both a single point of failure and a source of vibration. It allows the next exposure to begin almost immediately, optimizing frame rate and exposure time. It also avoids image smearing in the readout direction that would be present in even a frame-transfer CCD when the frame rate is high, and the exposure time approaches the transfer time.

However, since part of the surface of an interline CCD surface is non-responsive, and this device is front-illuminated, its Quantum Efficiency (QE) is low. To compensate, a grid of microlenses is incorporated into the CCD surface, which focuses incident light into the responsive regions. This raises the QE to 47% at 510 nm [TrueSense, 2013]. Unfortunately, this extra optical surface (and the electrode structure in a front-illuminated CCD) raises the potential for additional sources of detector Photo Response Non-Uniformity (PRNU) on sub-pixel scales: in some locations, light may be deflected toward the wrong pixel, or scattered and lost. This is of particular concern because the particular flowdown and budget allocation of science requirements in the BIT-STABLE mission means that the guide camera must enable position measurements with an accuracy better than 250 nm=0.045 pixels. To assess this camera's performance, we scan a small spot across the detector, recording its true position and comparing this to its measured position (and flux).

3.2. *Experimental setup*

To test the centroid determination accuracy, an illuminated pinhole was re-imaged onto the CCD surface. A piezo translation stage moved the pinhole source over the CCD surface with high positional accuracy. A separate confocal probe was used as a truth sensor, to independently monitor the stage position. Figures 7–9 show the experimental set-up.

The pinhole source was constructed using Thorlabs SM1 components. The interchangeable 5mm LED was held in a S1LEDM mount. A 5 mm BK7 ball lens, held in a customised SM1PL plug, focussed the light from the LED onto the pinhole (e.g. P5S). An SM1V10 focussing tube plus SM1L05 and SM1L10 lens tubes set the separation between the pinhole and the lens. A MAP107575-A matched achromat lens pair (75 mm focal length) with AR coating re-imaged the pinhole 1 : 1 onto the CCD

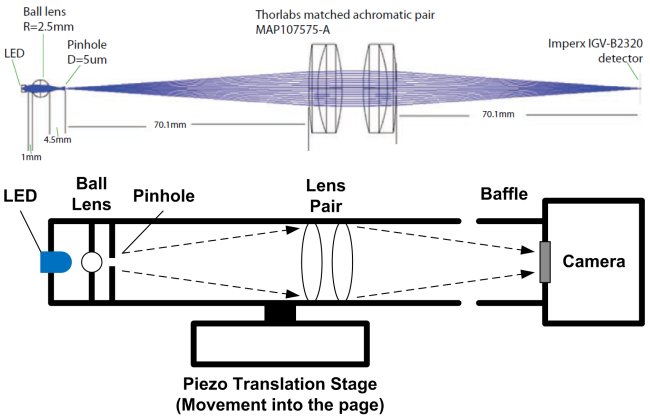

Fig. 7. Zemax model and schematic diagram of components used to shine a spot towards the camera, and move it around the detector at high precision.

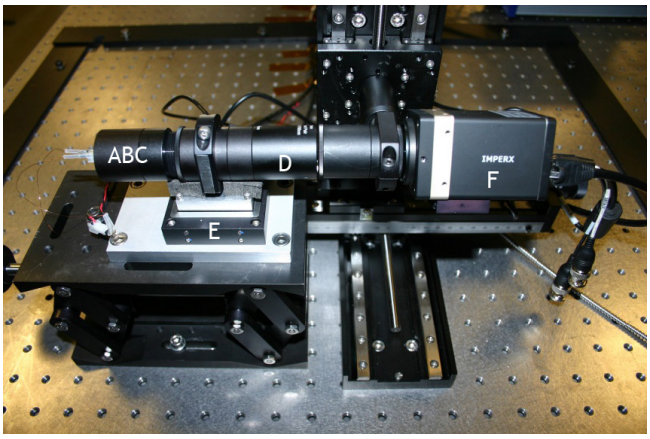

Fig. 8. Experimental set-up to scan a spot across the detector. Side view (as in figure 7). Key: A: 5mm LED; B: 5mm BK7 ball lens (not shown); C: mounted pinhole (not shown); D: 1:1 matched achromatic doublet pair; E: piezo translation stage; F: CCD camera with baffle tube. Confocal probe removed.

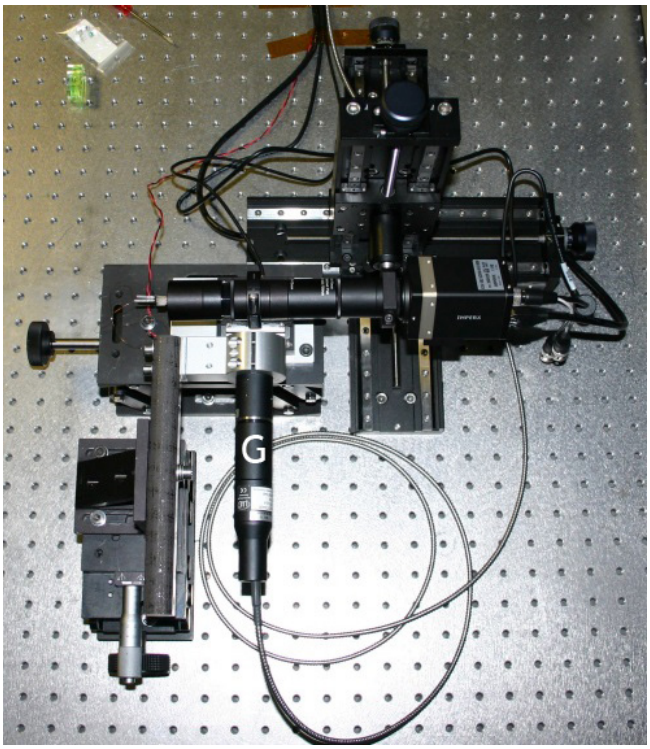

Fig. 9. Experimental set-up to scan a spot across the detector. Bird's eye view. Key: G: confocal probe.

surface. The lens pair has a quoted RMS spot size of 7.7 μm. The default setup used a 5 μm diameter pinhole and an LED with a central wavelength of 472 nm.

An LM371L 5 mA constant current source powered the LED. Thin enamelled copper wires were used to make the final connections to the LED to avoid restricting the piezo stage movement. A PICAXE 20X2 micro-controller and ULN2803A driver allowed the LED to be turned on and off via USB under the control of Python software running on a Windows PC. A Physik Instrumente (PI) P-625.1CD piezo translation stage (500 μm) was used to move the pinhole source over the CCD surface. The stage was operated in closed-loop with capacitive feedback. The 170 g mass of the pinhole source required the P, I, slew rate and notch filter parameters to be configured carefully to avoid causing oscillations in the stage position, and the optics were cushioned to damp any oscillations that did arise. The stage used for this experiment did not have digital dynamic linearization and so it was necessary to monitor the stage position with a separate Micro Epsilon confocal probe (300 μm) to allow the residual micron-level ripple in the stage linearity to be removed during the centroid calculation.

The Imperx camera and baffle tube were mounted on a manual XYZ translation stage to allow the spot position to be moved coarsely around the CCD surface, and to adjust the camera position for best focus at different LED wavelengths. The camera was rotated by $\sim$30° with respect to the piezo stage such that the image of the pinhole moved across the camera pixels in both X and Y simultaneously. Typically the pinhole source was moved by $5,000 \times 0.0333\,\mu$m or $20,000 \times 0.0111\,\mu$m steps and a light plus dark image pair collected at each position. This step size was chosen to be a non-integer fraction of the pixel size.

The GEV software delivered with the camera was used to capture images on a Windows PC. The software was configured for software trigger-

ing. Exposure times were adjusted to ensure a peak pixel count in the recorded images near 50% of full-well capacity, i.e. ∼2000 ADU (Analog to Digital Units) and avoid saturation. The Winger ultra-bright 5 mm LEDs used provided very high signal levels even at 5 mA drive current through a 5 μm pinhole. Consequently, exposure times were often as short as 3 μs. Additional software written in Python: turned the LED on or off via a USB port; stepped the position of the piezo stage via a serial port; collected measurements from the confocal probe via a second serial port; captured images via MOUSEEVENT function calls to press the software trigger button in the GEV camera software; logged the stage and probe positions to a `csv` file. Images were saved as 12-bit non-normalized TIFF files, using 128×128 pixel windows to minimise the amount of recorded data.

Data were collected with a 472 nm LED and various pinhole diameters (1μm, 2μm, 5μm, 15μm, 50μm), plus 516 nm, 624 nm and 850 nm LEDs with the 5μm diameter pinhole.

3.2.1. *Model and control data*

For each test run, we predicted the shape of the spot using Zemax geometrical optics (figure 7) to model a point source through the matched achromat lens pair at optimized focus, and convolving with a circular top hat of the pinhole diameter.

We also simulated a control set of mock data in which the detector had perfectly uniform response. To do this, we used software developed by [Rowe, Hirata & Rhodes, 2011] to analytically place an Airy disc spot (with an appropriate range of sizes) at 5,000 stepped locations with sub-pixel precision, before integrating it within pixels and realizing it with a flux of ∼2000 ADU in the peak pixel. We shall analyze this control data in the same way as our real data, to test the validity of our measurement methods in different regimes (it produces the solid lines in figures 12 and 13).

3.3. *Data analysis*

3.3.1. *Spot size measurement*

For large spots that are well resolved by the pixels, we measure their size directly from the images. To do this in an automated way on noisy data, we measure the full width at half maximum (FWHM) of the best-fit Gaussian, optimizing the center, flux and size via Levenberg-Marquardt minimisation. On the noiseless Zemax models, the FWHM of the best-fit Gaussian is typically within 4% of the FWHM of the full spot profile. After deconvolving from the pixel square, the measured FWHM is also within 6% of the Zemax model, for all but the largest pinholes (where the true shape of the spot is probably influenced by structure in the LED). For each test run, we record the mean FWHM measured in every exposure.

For small spots that are not well-resolved by the pixels, most of the light can fall within a single pixel. The best-fit Gaussian then just fits the size of a pixel, and direct measurement becomes impossible. For the test runs using a 472 nm LED and 1μm, 2μm, or 5μm pinholes, we use the FWHM of the best-fit Gaussian to the Zemax model. However, we indicate up to 10% uncertainty because of the difficulty focussing a spot that is not resolved by the detector and below the resolution limit of the lens.

3.3.2. *Spot position measurement*

To measure the observed position of spots, we adopt an iterative centroiding technique long-established for the high precision measurement of galaxy morphologies [Kaiser, Squires & Broadhurst, 1995; Rhodes, Refregier, & Groth, 2000]. We initially guess the center of the brightest pixel to be the center $\mathbf{x}_{\rm obs}^{i=0}$ of a spot image $I(\mathbf{x})$. To improve this guess, we compute the image's first moments, weighted by a Gaussian centered at the current guess, and iterate

$$\mathbf{x}_{\rm obs}^{i+1} = \frac{\iint \mathrm{e}^{-|\mathbf{x}-\mathbf{x}_{\rm obs}^{i}|^2/(2\sigma_w^2)}\ I(\mathbf{x})\ \mathbf{x}\ \mathrm{d}^2\mathbf{x}}{\iint \mathrm{e}^{-|\mathbf{x}-\mathbf{x}_{\rm obs}^{i}|^2/(2\sigma_w^2)}\ I(\mathbf{x})\ \mathrm{d}^2\mathbf{x}} \qquad (1)$$

where σ_w is a weight size (see below), and the integral is performed over all pixels within a radius of 5σ. For our images, the iteration converges to a precision better than 0.001 pixels within 2 (3, 4) steps for stars inside the central 1% (central 10%, all) of the pixel. In these tests, for which the calculations are not time-sensitive, we always perform 5 iterations and render this effect completely negligible.

To ensure a fair comparison, we fix the size of the weight function σ_w=2 pixels. The optimal choice of σ_w for fastest convergence depends upon the size and S/N of the spot, but we obtain consistent results for all but pathologically small values of σ_w. Note that $\sigma_w \to \infty$ recovers an unweighted centroid, but using this requires nonlinear thresholding to counteract the effect of noise [Rufino & Accardo,

2003; Li, Li & Liang, 2009].

We also record the centers of both circular and elliptical Gaussians fitted to the spot images. This is similarly accurate as the moment-based method on simulated resolved spots with high S/N, but is slow and less robust to noise. The mean offset between centroids calculated via moments and elliptical Gaussians in 20,000 real exposures is $(13.1\pm0.1, -4.6\pm0.1)\times10^{-3}$ pixels. Of most concern is that neither algorithm works for spots that are smaller than a pixel. Accurate fitting of very small spots requires precise knowledge of the spot wings [Tremsin *et al.*, 2003]. In principle, it could be possible to shift and coadd our images to super-resolve the spot, then measure its shape [Rowe, Hirata & Rhodes, 2011]. However, this situation will not be encountered in the real instrument; and since our data contains only a single spot, applying sub-pixel shifts in order to measure the sub-pixel position is circular logic. We therefore merely quantify the limitation of our algorithm, using simulated data where we know the true subpixel position.

3.3.3. *Spot flux measurement*

Our experimental setup has greatly simplified measurements of the total flux within a spot, even for unresolved spots. We sum the flux in all pixels whose centers lie within an aperture of radius 5 pixels or 2×FWHM (i.e. 4× the half width at half maximum) of the spot centroid after convolution with the pixel square, whichever is larger. Our final results, which compare average measurements from thousands of images, are robust within 2% for different choices of aperture radius.

3.3.4. *Input spot positions*

The input position of a spot $(x_{\rm true}, y_{\rm true})$ is set by the position of the translation stage. We assume that the translation stage moves the optics on a straight line, the distance along which is measured up to 20,000 times by the confocal probe as $\ell_{\rm confocal}(t)$. We then obtain the 2D input position

$$x_{\rm true}(t) = x_{\rm start} + g\ell_{\rm confocal}(t)\cos\theta + \ell_{\rm spline}(t)\cos\theta + x_{\rm spline}(t) \quad (2)$$

$$y_{\rm true}(t) = y_{\rm start} + g\ell_{\rm confocal}(t)\sin\theta + \ell_{\rm spline}(t)\sin\theta + y_{\rm spline}(t) \quad (3)$$

by fitting the following free parameters, assuming initially that $(x_{\rm true}, y_{\rm true}) \approx (x_{\rm obs}, y_{\rm obs})$:

- Starting subpixel position $(x_{\rm start}, y_{\rm start})$ of the first exposure.
- Angle θ between the spot scan line and the x-axis of the pixel grid.
- Pitch of the steps g between subsequent exposures, calibrating any linear gain in the confocal probe.

Plus, optionally:

- A cubic B-spline $\ell_{\rm spline}$, calibrating low spatial frequency, non-linear gain in the confocal probe or judders in the translation stage.
- Cubic B-splines $x_{\rm spline}$ and $y_{\rm spline}$ to perturb the true position of the spot at low spatial frequencies in the x and y pixel directions, accounting for residual vibration in the optical assembly, which needed to move, stop, and come to rest in between each exposure.

When we include the splines, we are careful to model variations only on scales larger than a pixel, to avoid spuriously introducing (or correcting) the kind of sub-pixel variations that we seek.

For a typical run with a 472 nm LED and 5 μm pinhole, the best-fit value for g is $1 + (4.0 \pm 0.2) \times 10^{-3}$. The best-fit $\ell_{\rm spline}$, $x_{\rm spline}$, and $y_{\rm spline}$ have an rms of 74 nm, 4.6 nm, and 6.7 nm respectively. The confocal probe had confirmed that our translation stage was juddering at around this level, so the fit is not surprising, but this demonstrates that our laboratory equipment is sufficiently stable to test the flight requirements. For the simulated data of a perfect detector (and experiment, with 472 nm LED and 5 μm pinhole), the best-fit splines have an rms of 1.2 nm, 0.9 nm and 0.7 nm, and are consistent with zero everywhere. This therefore demonstrates that the optional refinement is not introducing statistically significant perturbations.

3.3.5. *Modelling PRNU variation in 2D*

We have not scanned a spot in a 2D pattern around one pixel, but in a 1D pattern across several pixels. To map 2D PRNU, we assume that any pattern is repeated identically in all the pixels within small ($\sim 25\times15$ pixel) patches of the detector. We then wrap the sub-pixel input spot positions such that $0 < x_{\rm true} < 1$ and $0 < y_{\rm true} < 1$, mapping them all back onto a virtual pixel that becomes well sampled (figure 10).

To check the validity of this assumption, we have: (i) performed the test runs (including some repeat runs with otherwise identical configuration)

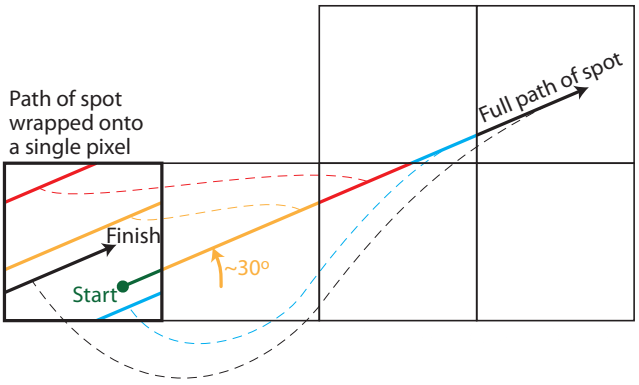

Fig. 10. We scan a spot in a continuous 1D line across several pixels. As justified in the text, we then assume that sub-pixel PRNU is identical within small patches of the detector. We then map the input spot positions back onto a single virtual pixel, so the entire 2D pattern of PRNU is well sampled.

Table 1. Sizes of the spot produced with various experimental configurations and scanned across the detector. Measurements from a Zemax model or real images are consistent within 6% (see Section 3.3.1).

LED wavelength [nm]	Pinhole diameter [μm]	Spot FWHM [μm]
472 (blue)	1	3.57
472	2	3.86
472	5	5.05
516 (green)	5	7.02
624 (red)	5	12.44
472	15	14.35
850 (near-IR)	5	17.06
472	50	42.74

in different parts of the detector; the results from these runs are consistent with each other, within experimental errors. (ii) explicitly tried to fit patterns across 2×2 groups of pixels. It might be possible to have such patterns if e.g. a chip mask used to manufacture devices with an RGB Bayer mosaic were adapted to create this monochromatic device. We find that power in a 2-pixel pattern is less than 1% of that in the 1-pixel pattern, and consistent with zero. We therefore assume that the tiny residuals on these spatial scales are due to vibrations in the translation stage, and allow them to be fitted away by the B-splines in equations (2) and (3).

3.4. *Results*

3.4.1. *Spot sizes*

The sizes of the spot produced by different combinations of LED wavelength and pinhole size are listed in table 1. Colors indicative of the LED wavelengths are used for the data points in figures 12 and 13.

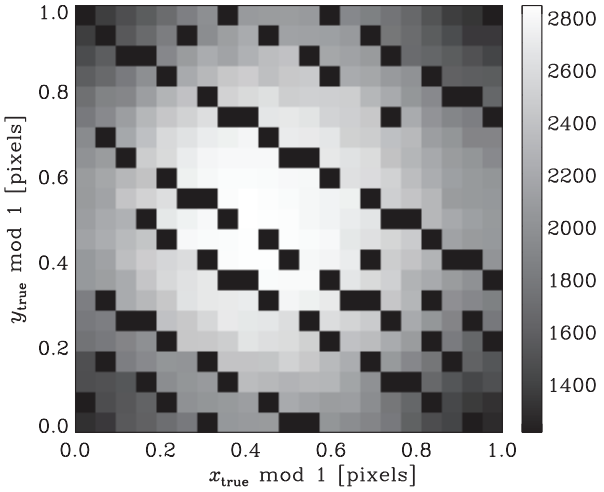

Fig. 11. Pixel Response Non-Uniformity: the total flux recorded (within a 5 pixel aperture) when our smallest spots (1 μm pinhole, 472 nm LED) are shone at different sub-pixel positions. Black regions contain no data, and the scanning strategy of figure 10 is apparent. The response of a perfect detector would be uniform, but here some flux is lost near the edges of pixels.

3.4.2. *Flux loss as a function of sub-pixel position*

For very small spots, the variation in observed total flux at different sub-pixel positions $\Delta F(x_{\rm true}, y_{\rm true})$ is shown in figure 11 and well fit by

$$\Delta F = F_{\rm peak}\big[1 - A_x \sin^4(\pi(x_{\rm true} - x_c)/1\,\text{pixel}) - A_y \sin^4(\pi(y_{\rm true} - y_c)/1\,\text{pixel})\big]. \quad (4)$$

This pattern peaks at the center of the pixel, (x_c, y_c), and drops at the sides in a continuous way between adjacent pixels; it thereby mirrors the likely shape of the microlens array focussing light onto the responsive areas of the CCD. The exponent 4 represents a compromise between values that best fit data from test runs with various spot sizes; this varies between 3–5 if left as a free parameter, but we fix it to simplify comparison between runs.

For test runs with every size spot, we determine best-fit values for the free parameters $F_{\rm peak}$, (x_c, y_c), and A_x, A_y using Levenberg-Marquardt minimisation. Measured values for A_x and A_y are shown in figure 12; a perfect detector would have $A_x = A_y = 0$, as indeed recovered in our mock test data with simulated images.

The flux loss detected for small spots is therefore real and significant. The greater flux loss at the sides of a pixel than at the top/bottom ($A_x > A_y$) is also expected, due to the asymmetric detector architecture in which the responsive photodi-

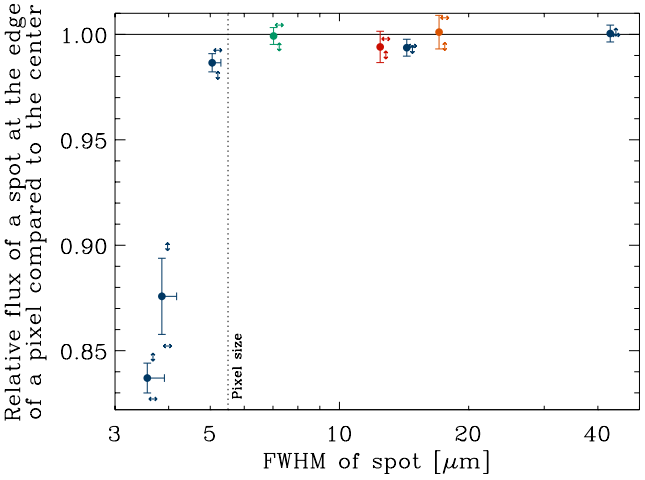

Fig. 12. Pixel Response Non-Uniformity: the fraction of flux lost at the edges of a pixel, from an input spot of varying widths. Data points show $1-(A_x+A_y)/2$, which are defined in equation (4). Offset slightly to the right and without error bars for clarity, the left-right arrows show $1-A_x$ and the up-down arrows show $1-A_y$ individually. The leftmost data corresponds to figure 11. A perfect detector would have $A_x = A_y = 0$.

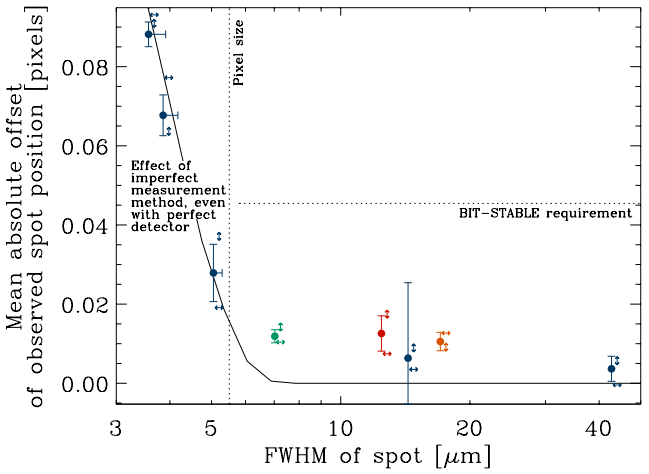

Fig. 13. Pixel Response Non-Uniformity: the spurious shift in the position of an input spot of varying widths at the edges of a pixel. Following equation (4), data points show $(B_x + B_y)/2$. Offset slightly to the right and without error bars for clarity, the left-right arrows show B_x and the up-down arrows show B_y individually. A perfect detector would have $B_x = B_y = 0$, and BIT-STABLE requires performance below the horizontal line.

ode and non-responsive interline readout columns lie to the sides of the pixels. Indeed, the best-fit value for (x_c, y_c) in figure 11 is offset laterally to $(0.468 \pm 0.003, 0.502 \pm 0.003)$ pixels, although this result varies slightly with position within the detector, whose microlenses are positioned in a telecentric pattern. Our results thus indicate that some light is still lost into the nonresponsive regions, despite the microlens array.

However, the observed flux in large spots that are resolved by the pixels (especially if they are Nyquist sampled or better) are unaffected by this PRNU within our experimental precision. Some flux loss may indeed contribute to the overall QE (here degenerate with $F_{\rm peak}$) but, with the loss averaged over many locations, the flux response appears uniform for a spot positioned anywhere.

3.4.3. *Centroid shift as a function of sub-pixel position*

The centroid of a spot can be shifted both by a gradient in detector sensitivity across the spot, or (more likely) by optical distortion through the lenslet array. Once again reflecting the lens geometry, we fit offsets in the observed centroid compared to its input position $(\Delta x, \Delta y) \equiv (x_{\rm obs} - x_{\rm true}, y_{\rm obs} - y_{\rm true})$ as it moves within a pixel as

$$\Delta x = B_x \sin\left(2\pi(x_c - x_{\rm true})/1\,\text{pixel}\right) \quad (5)$$

$$\Delta y = B_y \sin\left(2\pi(y_c - y_{\rm true})/1\,\text{pixel}\right). \quad (6)$$

This pattern is zero in the middle of the pixel and at the edges, where symmetry also ensures that the lenslet must also be flat. Nuisance parameters (x_c, y_c) are degenerate with $(x_{\rm start}, y_{\rm start})$ in equations (2) and (3), but are left as free parameters here for completeness. Positive values of B_x or B_y indicate a shift towards the center of a pixel. This function, including the (lack of) exponent, is a very good fit to the behavior of an imperfect position measurement algorithm when applied to simulated images from a perfect detector (see §3.3.2). To try to dig out any other signal, we have also tried varying the exponent, or doubling the frequency of the sinusoidal variations. However, no other significant signal is found. For example, for the test run using the 472 nm LED and 5 μm pinhole, the best-fit exponent is 1.01 ± 0.17, the amplitudes B_x, B_y are unaffected within 1% by the addition of extra terms with double frequency, whose best-fit amplitudes are less than 3% of B_x and B_y.

The mean absolute offsets, $\sqrt{\Delta x^2 + \Delta y^2}$, averaged within an entire pixel, are shown for different sized spots in figure 13. The difficulty of position measurements for small spots makes it hard to distinguish between their expected centroid shifts

and imperfect algorithmic behavior with any statistically precision. For a Nyquist-sampled spot, the mean offset is around 0.01 pixels=55 nm. For a large spot, the mean offset appears to tend to zero, as expected if the individual shifts within many pixels average away. The mean offset is the appropriate quantity to compare to requirements because noise in the position measurement from an individual image with a guide camera will be dominated by photon-counting shot noise. The *peak* offset anywhere within the pixel is ~ 1.5 times greater than the mean offset but even this, for a Nyquist-sampled PSF, is still within requirements for BIT-STABLE.

4. Windowed readout

4.1. *The problem*

The BIT-STABLE mission requires $>100\times100$ pixel images of a guide star at > 50 Hz. Because of the finite readout speed and data transfer rate over 100Mbps ethernet, it will not be possible during flight to read out the full frame and crop the region of interest in postprocessing. During further camera testing at the Jet Propulsion Laboratories, California Institute of Technology (JPL), we noticed artifacts in images obtained using Windowed Readout Mode. Compared to images obtained using Full Frame Readout Mode, we observed an apparent overall loss of sensitivity, and a spatially varying sensitivity gradient in the readout direction.

Any loss of CCD sensitivity will degrade photometric measurements, and raise the magnitude limit of guide stars for the BIT-STABLE mission. Uncorrected spatial variations in sensitivity will also affect the camera's measurements of (relative) position. The effect can be computed by inserting a multiplicative factor $S(\mathbf{x})$ into the numerator and denominator of equation (1). Let us assume that the PSF image $I(\mathbf{x})$ is a Gaussian of width σ_p, centered at $\mathbf{x}_{\rm true} = \mathbf{x}_{\rm o} + \delta\mathbf{x}$, where $\mathbf{x}_{\rm o}$ is a nearby reference point. For a sensitivity variation in the readout direction expressed as a Taylor series about $\mathbf{x}_{\rm o}$

$$S(\mathbf{x}) = a + b(y - y_{\rm o}) + c(y - y_{\rm o})^2, \qquad (7)$$

the observed y position of the PSF image becomes

$$y_{\rm obs} = y_{\rm true} + \frac{\sigma^2(b + 2c\delta y)}{a + b\delta y + c\delta y^2 + c\sigma^2} \qquad (8)$$

$$\approx y_{\rm true} + \frac{b}{a}\sigma^2 + \frac{2c}{a}\sigma^2\delta y \qquad (9)$$

to first order in δy and leading order in b/a and c/a, and where $\sigma^2 = \sigma_p^2 + \sigma_w^2$.

A constant sensitivity gradient would thus shift all measured positions by a constant offset – which would be unnoticeable in either guide star imaging (where only the relative position between exposures is measured) or our spot projection tests (where we fitted the absolute position $\mathbf{x}_{\rm start}$ as a free parameter). Only second order (and higher) spatial derivatives of sensitivity affect measurements of the offset between two sources, or a moving source observed in successive exposures, via

$$\frac{{\rm d}y_{\rm obs}}{{\rm d}y_{\rm true}} = 1 + \frac{2c}{a}\sigma^2. \qquad (10)$$

4.2. *Experimental setup*

This testing was performed at the Jet Propulsion Laboratories (JPL). The camera was exposed to a near-flat DC light source, and images were captured in both Full Frame and Windowed Readout Mode (under the same lighting conditions, camera modes, and exposure times).

4.3. *Results*

4.3.1. *Observed behavior*

In Full Frame Readout Mode, the received flat field image is roughly flat, as expected. Figure 14 shows pixels (500:700, 500:700) cropped from full frame images, averaging over 35 consecutive exposures to reduce the amount of readout and shot noise[1]. However, the same 200×200 pixels captured in Windowed Mode exhibit a reduced signal and a "waterfall"-like gradient (figure 15). To characterize the drop in sensitivity, we divide the Windowed Mode image (figure 15) by the corresponding pixels of the Full Frame Readout Mode image (figure 14). Spatial variation in the result is well fit by

$$S(\mathbf{x}) = (0.841 \pm 0.001) - (0.253 \pm 0.001)\exp\left(\frac{500 - y}{28.8 \pm 0.18}\right). \qquad (11)$$

The 84.1% drop in sensitivity is constant within a few percent if the window is moved around the CCD. The drop is negligible for the largest possible (e.g. 1700×2000 pixel) windows, and becomes progressively worse as the window size is reduced.

[1] By considering flat field images with different exposure times, we find that the readout noise and dark current are consistent with the $12\,{\rm e}^-$ rms and $1\,{\rm e}^-/{\rm s}$ values quoted in the detector specifications [TrueSense, 2013].

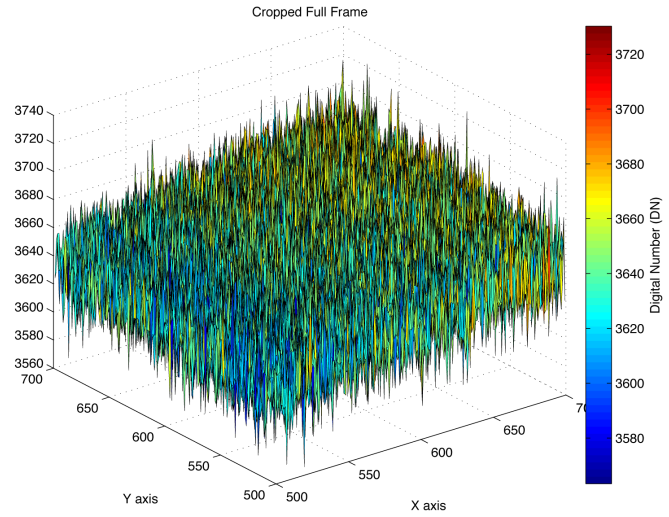

Fig. 14. Flat field image obtained in Full Frame Readout Mode. A cropped, 200×200 pixel region is shown. To reduce readout and shot noise, the image has been averaged over 35 consecutive exposures.

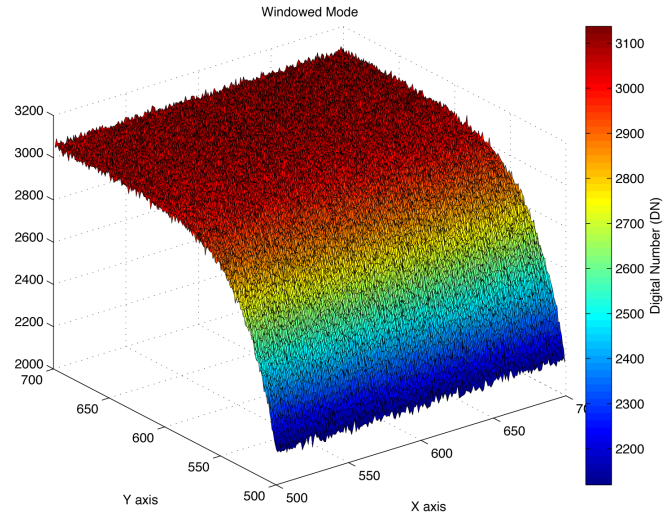

Fig. 15. Flat field image obtained in Windowed Mode. The 200×200 pixel window corresponds to the same pixels as those shown in figure 14. To reduce readout and shot noise, the image has been averaged over 35 consecutive exposures.

Dark frames observed in Windowed Readout Mode also exhibit a spurious excess signal. An additive bias of ~ 1 ADU is present in the first row, decreasing linearly in the readout direction to ~ 0.1 ADU after 50 rows. The spurious signal alternates by $\pm 10\%$ on odd/even columns. However, it is constant, regardless of the exposure time.

4.3.2. *Possible explanation*

The camera manufacturer reported that the sensor does not have a hardware fast line dump, which

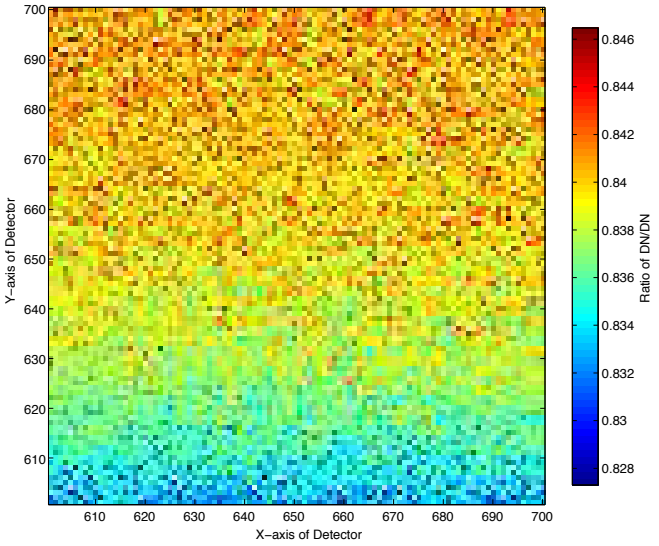

Fig. 16. The net effect on sensitivity, when using our workaround of obtaining a larger window than required. This shows the ratio of counts in the Windowed Mode image compared to those in the Full Frame Readout Mode image, in the 'flat' 100×100 pixel region at the top of the window.

is used in other devices to quickly discard charge from all pixels in a row and rapidly advance readout toward a restricted region of interest. To produce windowed readout from the IGV-B2320-M, this feature is being emulated. The exponential form of equation (11) is reminiscent of charge transfer problems [e.g. Massey *et al.*, 2010; Short *et al.*, 2013], perhaps because the clock speed is being pushed higher than the 40 MHz maximum specification [TrueSense, 2013]. The loss of sensitivity may then be caused by a buildup of (undumped) charge, which affects the sensitivity of the output amplifier. It would be useful to compare this camera's performance with that of the IGV-B2020-M, which utilises the Truesense KAI-04022 detector (2048×2048 $7.4\,\mu$m pixels) and which does have a fast line dump.

Note that if the cause of sensitivity loss is indeed undumped charge, a flat field with a high level of charge outside the window presents a worst case scenario. Images from a guide camera that contain mostly blank sky should be less affected. It would be interesting to make photometric measurements of an isolated spot of constant flux as it moves a long distance across the detector in Windowed Readout and Full Frame Readout Modes. However, our spot projector lab is not currently available for further tests of this camera.

4.3.3. *Mitigation strategy*

The low level of spurious additive signal can be removed via subtraction of dark frames obtained in flight. Multiplicative sensitivity variation could be reduced via a flat field. However, in case the effect is not constant, or signal-dependent, we require a more robust mitigation strategy.

If the first 100×100 pixel region to be read out were adopted, without any correction, as the region of interest (as it was in our spot projection tests of section 3), it would have a mean sensitivity of 0.770 ± 0.003 and a steep local gradient in sensitivity. Approximating equation (11) as a Taylor series near the middle of this region (i.e. with $y_o = 550$) yields $a = 0.796$, $b = 1.55 \times 10^{-3}\,\mathrm{pixel}^{-1}$ and $c = -5.37 \times 10^{-5}\,\mathrm{pixel}^{-2}$. With a Nyquist sampled PSF, for which $\sigma_p = 2/2.3548$ pixels and hence $\sigma = 2.2$ pixels, equation (9) suggests that $|\mathrm{d}y_{\mathrm{obs}}/\mathrm{d}y_{\mathrm{true}} - 1| \approx 6.4 \times 10^{-4} = 3.5\,\mathrm{nm/pixel}$. This would formally fail mission requirements if the spot frequently moved by 100 pixels between successive exposures (but the closed-loop guidance would converge so long as typical movements are smaller). Across the ~ 15 pixel rows traversed in our spot projection tests, this large-scale variation could explain a fraction of the 74 nm rms fluctutation of ℓ_{spline} that we had previously ascribed to juddering of the translation stage.

As a workaround to further militate against sensitivity loss and its spatial variation, we capture a taller window than is necessary, then discard the first 100 rows and retain only the region farthest from the readout register – in which the sensitivity is higher and approximately flat (figure 16). For example, if the region of interest is pixels (600:700, 600:700), we capture pixels (600:700, 500:700). At 50 Hz frame rate, this 100×200 pixel window remains well within the maximum ethernet transfer rate (see section 2.2.2) and the camera's maximum frame rate (which depends only on the number of rows read out). In the retained region, the mean sensitivity is 0.839 ± 0.002 and its spatial variation is well fit by equation (7) with $y_o = 650$, $a = 0.839 \pm 0.002$, $b = (6.88 \pm 5.36) \times 10^{-5}$ pixel^{-1}, and $c = (-1.04 \pm 2.08) \times 10^{-6}\,\mathrm{pixel}^{-2}$. Again assuming a Nyquist sampled PSF, this indicates that $|\mathrm{d}y_{\mathrm{obs}}/\mathrm{d}y_{\mathrm{true}} - 1| \approx 1.2 \times 10^{-5} = 0.064\,\mathrm{nm/pixel}$. This level of error will have a negligible effect on the device's performance as a fine guidance camera, so long as the input motion of the guide star between exposures keeps it within the 100 pixel region of interest — and the mission would fail anyway if it moved farther than that.

5. Conclusions

We have successfully demonstrated that an unmodified commercial off-the-shelf front-illuminated progressive scan interline transfer CCD guide camera can meet the requirements of the closed-loop star-guiding and tip-tilt correction system for the BIT-STABLE balloon mission.

The camera can withstand the $+70°\mathrm{C}$ / $-40°\mathrm{C}$ survivability and $+60°\mathrm{C}$ / $-30°\mathrm{C}$ operating temperature requirements in near-vacuum (pressures of 4mBar or below). Incident light has to pass through a microlens array and be focussed onto photoresponsive regions of silicon in order to be detected. Some light remains undetected, particularly at the sides of pixels adjacent to the interline readout columns. However, this effect is averaged out over several pixels for a Nyquist-sampled spot, whose centroid can be measured to 1% (1.5%) of a pixel in an average (worst) case sub-pixel position. artifacts in sensitivity and flat field are created in the first ~ 100 rows of a windowed image because the detector does not have a hardware fast dump. The required performance can be recovered by capturing larger images, and discarding the first 100 rows.

Data transfer from the camera works reliably in 100 Mbps Full Duplex ethernet mode, but not in GigE ethernet mode. This lowers the theoretical maximum frame rate for full-frame images, but remains within the requirements of BIT-STABLE.

Acknowledgments

The authors thank Barnaby Rowe for software to create pixellated PSF models, Carl Christian Liebe for guidance on noise analysis, and Roger Smith for discussions of charge transport. We acknowledge support from Durham University's Centre for Advanced Instrumentation, Edinburgh University sustainability fund, UK Science and Technology Facilities Council (grant number ST/L00075X/1), European Research Council (grant number MIRG-CT-208994), and the Leverhulme Trust (grant number PLP-2011-003). This work was supported in part by the National Science Foundation under Grant No. PHYS-1066293 and the hospitality of the Aspen Center for Physics. RM is supported by a a Royal Society Research Fellowship. The work of HC, LLJ, MM and SS was carried out at the Jet Propulsion Laboratory, California Institute of Technology,

under contract with the National Aeronautics and Space Administration (NASA).

References

Kaiser N., Squires G. & Broadhurst T., 1995, ApJ, 449, 460

Li Y., Li D. & Liang Y., 2009, Proc. SPIE 7157

Massey R., Stoughton C., Leauthaud A., Rhodes J., Koekemoer A., Ellis R. & Shaghoulian E., 2010, MNRAS 401, 371

Pascale E. *et al.* 2008, ApJ 681, 400

Rhodes, J., Refregier, A. & Groth, E., 2000, ApJ, 536, 79

Rowe B., Hirata C. & Rhodes J., 2011, ApJ 741, 46

Rufino G. & Accardo D., 2003, Acta Astronautica 53 (2), 135

Short A., Crowley C., de Buijne J. & Prodhomme T., 2013, MNRAS 430, 3078

Tremsin A., Vallerga J., Siegmund O. & Hull J., 2003, Proc. SPIE 5164

Truch M. *et al.* 2009, ApJ 707, 1723

Truesense KAI-04050 image sensor device performance specification document, revision 4.0 PS-0009, 2013 October 15 http://www.truesenseimaging.com/products/interline-transfer-ccd/35-KAI-04050

U.S. Standard Atmosphere, 1976, NASA-TM-X-74335, NOAA-S/T 76-1562